

PATENT
PD-00-1014

**METHOD OF STORING BIOS MODULES AND
TRANSFERRING THEM TO MEMORY FOR EXECUTION**

Curtis E. Stevens

METHOD OF STORING BIOS MODULES AND TRANSFERRING THEM TO MEMORY FOR EXECUTION

BACKGROUND

The present invention relates generally to computer systems and methods, and more particularly, to a method that may be used to store basic input output system (BIOS) modules of a computer in a protected area, transfer the BIOS modules to
5 memory, and execute the BIOS modules to configure the computer.

A BIOS, or basic input output system, refers to a firmware program that is typically stored in nonvolatile memory (flash memory), and which brings up (initializes) a computer system when it is powered on. In conventional computer systems, the BIOS is typically loaded as a single block of code into random access memory (RAM) where
10 it is accessed by the CPU of the computer system to operate the system.

Prior art hard disk drives are capable of safely storing firmware (BIOS) in a protected area on spinning media in a tightly integrated system. A tightly integrated system is one where components used for basic operation of the hard drive are interdependent. When the hard drive is a component of a larger system, typically no
15 other firmware components can be safely stored on the hard drive. The protected area is referred to as a vendor protected area. The vendor protected area is typically reserved for firmware of the hard disk drive manufacturer. In general, a fixed area of less than one megabyte of hard disk space is reserved for the vendor protected area.

With regard to prior art specifically related to the present invention, US Patent
20 No. 5,835,760 issued November 10, 1998, entitled "Method and arrangement for providing BIOS to a host computer", discloses operation of a computer system having a

host computer with system RAM. The host computer uses a BIOS to control the operation of the computer system, and requires the host computer to obtain the BIOS and store it in the system RAM. A first portion of the BIOS (initialization code) is stored in a read only memory (ROM) which is not a flash memory, and a second
5 portion of the BIOS is stored in a protected area of a mass memory storage peripheral computer device (hard disk) rather than in ROM. It is stated that the second portion of the BIOS may be expansion BIOS associated with a particular peripheral computer device or system BIOS associated with the host computer.

10 In operation, as is stated in US Patent No. 5,835,760, the first portion of the BIOS (initialization code) is read out of the ROM into the host computer. The host computer executes instructions contained in the first portion of the BIOS and reads the second portion of the BIOS (i.e., the entire BIOS) from the protected area of the hard disk drive into the system RAM. The second portion of the BIOS is executed to operate peripheral computer devices or the computer.

15 US Patent No. 5,835,760 does not disclose or suggest the use of a flash memory or nonvolatile memory, and does not disclose or suggest storing any portion of the BIOS in a flash memory. Furthermore, US Patent No. 5,835,760 does not disclose or suggest selectively transferring modules of the BIOS into RAM, only the entire BIOS associated with the peripheral computer devices or the computer.

20 It is an objective of the present invention to provide for improved methods that may be used to store basic input output system (BIOS) modules of a computer in a protected area, transfer the BIOS modules to memory, and execute the BIOS modules to configure the computer. The present invention is an improvement over the teachings of US Patent No. 5,835,760.

25

SUMMARY OF THE INVENTION

To accomplish the above and other objectives, the present invention provides for methods that process basic input output system (BIOS) modules of a computer to initialize a computer. The methods store basic input output system (BIOS) modules of
30 the computer in a protected area, transfer the BIOS modules to system memory, and execute the BIOS modules to configure the computer. The computer includes a system memory, a central processing unit (CPU), a computer bus, a critical nonvolatile storage device and optional secondary nonvolatile storage device(s). Selected BIOS modules required for operation of the computer are transferred from the critical nonvolatile
35 storage device, and optionally the protected area of the secondary nonvolatile storage device, to system memory, and are executed to initialize the computer.

The critical nonvolatile storage device may be flash memory, read only memory (ROM), programmable read only memory (PROM), erasable programmable read only memory (EPROM), or electrically erasable programmable read only memory (EEPROM), or other device or technology that the CPU may use to execute an initial set of instructions. The secondary nonvolatile storage device may be a hard disk drive, a compact disk (CD) drive, a digital video disk (DVD) drive, a floppy disk drive, a Zip drive, a SuperDisk drive, a Magneto-Optical disk drive, a Jazz drive, a high density floppy disk (HiFD) drive, flash memory, read only memory (ROM), programmable read only memory (PROM), erasable programmable read only memory (EPROM), or electrically erasable programmable read only memory (EEPROM), or any other device capable of preserving data in the event of a power-off condition.

The basic input output system (BIOS) is firmware that is typically stored in solid state non-volatile memory. The present invention allows multiple CPUs and system devices to be initialized from firmware (BIOS) stored on non-solid state system resources. The present invention allows the system to safely store and reliably retrieve firmware from a non-solid state device.

In one embodiment of the present invention, a minimal amount of BIOS initialization code is stored in a first portion of the critical nonvolatile storage device. The minimal initialization code is operative to initialize the CPU and the system memory. Additional BIOS modules are stored in the critical nonvolatile storage device or a protected area of the secondary nonvolatile storage device. A dispatch manager is stored in a second portion of the critical nonvolatile storage device. The dispatch manager is operative to selectively load and iteratively execute a predetermined number of tasks relating to complete initialization of the computer.

After the computer is turned on, the minimal initialization code is executed to initialize the CPU and the system memory. The dispatch manager is copied from the critical nonvolatile storage device to the system memory. The dispatch manager sequentially executes the predetermined number of tasks to initialize the computer.

The plurality of BIOS modules that are operative to control operation of the computer may be stored in a protected area of the critical nonvolatile storage device. The dispatch manager executes to sequentially copy selected BIOS modules identified in its list of tasks to execute. These BIOS modules are copied from the protected area of the secondary nonvolatile storage device to the system memory and are then executed.

The dispatch manager may also operate to determine which BIOS modules are required for operation of the computer and determine if a required BIOS module is stored in the system memory. If the required BIOS module is stored in memory, then the dispatch manager executes the stored BIOS module. However, if the required

BIOS module is not in memory, then the dispatch manager copies the required BIOS module from the secondary nonvolatile storage device to the system memory and then executes the copied BIOS module. The dispatch manager repeats this procedure until all required BIOS modules are copied and executed. Once all of the required BIOS modules are copied into the system memory and executed, the operating system of the computer is launched and normal computer operation ensues.

The present invention thus allows system firmware (BIOS) to be stored on nonvolatile storage devices that are loosely integrated with the system. The basis of this operation is the ability to create a protected area on the nonvolatile storage device. In a hard disk environment, this protected area is distinct from the vendor-protected area discussed previously. This protected area can be accessed by the system firmware (the minimal initialization BIOS code), and BIOS modules required for the system are read from the protected area into system memory. The nonvolatile storage device denies all other access to the protected area allocated to the BIOS.

Heretofore, no computer or BIOS manufacturer has stored initialization BIOS in a first portion of a nonvolatile storage device which is executed to load a dispatch manager from a second portion of the nonvolatile storage device into system memory which is subsequently used to copy individual BIOS modules stored in a protected area and execute the copied individual BIOS modules to initialize the computer.

As was mentioned above, conventional hard disk drives include a hidden vendor-protected area where hard drive firmware is stored. However, in another embodiment of the present invention, the secondary nonvolatile storage device, and in particular a hard disk drive, is configured to have a second hidden protected area where the BIOS modules are stored. This BIOS module storage area is distinct from the vendor-protected area normally found on hard disk drives.

BRIEF DESCRIPTION OF THE DRAWINGS

The various features and advantages of the present invention may be more readily understood with reference to the following detailed description taken in conjunction with the accompanying drawing, wherein like reference numerals designate like structural elements, and in which:

Fig. 1 illustrates an exemplary prior art computer system that implements conventional BIOS loading and execution;

Fig. 2 illustrates a computer system that implements an exemplary method in accordance with the principles of the present invention that is used to store basic input output system (BIOS) modules of a computer in a protected area and transfer selected BIOS modules to system memory;

Fig. 3 illustrates details of the computer system shown in Fig. 2;

Fig. 4 illustrates a layout of an exemplary hard disk drive implemented in accordance with the principles of the present invention;

Fig. 5 is a flow diagram illustrating one exemplary method in accordance with the principles of the present invention that may be used to protect BIOS modules in a secondary nonvolatile storage device; and

Fig. 6 is a flow diagram that illustrates another exemplary method in accordance with the principles of the present invention that stores basic input output system (BIOS) modules of a computer in a protected area, transfers selected BIOS modules to system memory and executes the modules.

DETAILED DESCRIPTION

Referring to the drawing figures, Fig. 1 illustrates a simplified block diagram of an exemplary prior art computer system 10 that implements conventional BIOS loading and execution. This computer 10 corresponds to the computer system disclosed in US Patent No. 5,835,760.

This computer 10 has a central processing unit (CPU) 11 that is coupled to a read only memory (ROM) 12. The CPU 11 is coupled to a system memory 13, such as a random access memory (RAM) 13. The CPU 11 is also coupled to a hard disk drive 20 by way of a system bus 14, such as a Peripheral Component Interconnect (PCI) bus 14, for example. The ROM 12 stores initialization code that initializes the CPU 11, the system memory 13, the system bus 14, and the hard disk drive 20. A single piece of monolithic code comprising the basic input output system (BIOS) is stored in a protected area of the hard disk drive 20. This protected area is the vendor protected area discussed above.

A first portion of the BIOS (the initialization code) is stored in the ROM 12. A second portion of the BIOS (which a single monolithic block of code) is stored in a protected area of the hard drive disk 20. In operation, the initialization code is read out of the ROM 12 by the CPU 11. The CPU 11 executes instructions contained in the initialization code and reads the second portion of the BIOS (the entire BIOS) from the protected area of the hard disk drive 20 into the system RAM 13. The second portion of the BIOS is executed to operate peripheral computer devices or the computer 10.

Referring to Fig. 2, it illustrates a computer 10a that implements an exemplary method 50 (Fig. 6) in accordance with the principles of the present invention. The method 50 is used to store basic input output system (BIOS) modules of the computer 10a (Fig 3) in a protected area, transfer selected BIOS modules to system memory 13, and execute the selected BIOS modules.

The computer 10a comprises a central processing unit (CPU) 11 that is coupled to a critical nonvolatile storage device 15. The critical nonvolatile storage device 15 may be flash memory, a read only memory (ROM), a programmable read only memory (PROM), an erasable programmable read only memory (EPROM), an electrically erasable programmable read only memory (EEPROM), or other device or technology that the CPU 11 can use to execute an initial set of instructions.

The CPU 11 is also coupled to a system memory 13, such as a random access memory 13. The CPU 11 may be coupled to a secondary nonvolatile storage device 20a by way of a system bus 14, such as a Peripheral Component Interconnect (PCI) bus 14, for example. The secondary nonvolatile storage device 20a may be a hard disk drive, a compact disk (CD) drive, a digital video disk (DVD) drive, a floppy disk drive, a Zip drive, a SuperDisk drive, a Magneto-Optical disk drive, a Jazz drive, a high density floppy disk (HiFD) drive, flash memory, read only memory (ROM), programmable read only memory (PROM), erasable programmable read only memory (EPROM), electrically erasable programmable read only memory (EEPROM), or any other device or technology capable of preserving data in the event of a power-off condition.

A first portion of the critical nonvolatile storage device 15 stores initialization code 16 that is operative to initialize the CPU 11 and the system memory 13. A second portion of the critical nonvolatile storage device 15 stores a dispatch manager 17 that contains a list of tasks, which must execute to fully initialize the computer 10a. The dispatch manager 17 is operative to selectively load and iteratively execute a number of tasks relating to complete initialization of the computer.

In operation, when the computer 10a is turned on, the initialization code 16 is run to initialize the CPU 11 and the system memory 13. The dispatch manager 17 is then loaded into the system memory 13. The dispatch manager 17 executes the list of tasks contained therein to cause all required BIOS modules to be loaded into the system memory 13 and must be executed.

The dispatch manager 17 determines whether each required BIOS module in the system memory 13, and if it is not, finds, loads and executes each required BIOS module. The BIOS modules may be located in the critical nonvolatile storage device 15 (flash memory) or in the secondary nonvolatile storage device 20a, including any of the critical or secondary nonvolatile storage devices 20a identified above.

Fig. 3 illustrates details of the computer system 10a shown in Fig. 2. The computer system 10a is shown comprising the central processing unit (CPU) 11, the flash memory 12, and the system memory 13. The CPU 11 is coupled to the secondary nonvolatile storage device 20a by way of a system bus 14, such as a Peripheral

Component Interconnect (PCI) bus 14, for example. A video controller 18 is coupled to the system bus 14 and to the system memory 13. A host controller 19 is coupled between the PCI bus 14 and to the secondary nonvolatile storage device 20a. A secondary nonvolatile storage device 20a implemented as a hard disk drive 20a may be an IBM DJNA-371350 ATA hard disk drive 20, for example. The hard disk drive 20a comprises a spindle 21, which is coupled to a central processing unit (CPU) 22 having a volatile memory 23.

Fig. 4 illustrates a layout of an exemplary hard disk drive 20a implemented in accordance with the principles of the present invention. The media of the hard disk drive 20a is broken up into three distinct areas. The first is a vendor protected area 25. The vendor protected area 25 is typically reserved for firmware of the manufacturer of the hard disk drive 20a. In general, a fixed area of less than one megabyte of hard disk space is reserved for the vendor protected area 25.

In accordance with the present invention, a second protected area 27 of the hard disk drive 20 (the secondary nonvolatile storage device 20a), which may be referred to as a BIOS protected area 27, contains a plurality of individual BIOS modules for the computer 10a. The BIOS protected area 27 may be created on the hard disk drive 20a using NCITS D1367 (PARTIES), NCITS D1321 (ATA/ATAPI-5), and NCITS D996 (SCSI-3 Block Commands) specifications.

The PARTIES specification specifies how to organize data on the secondary nonvolatile storage device 20a. The ATA/ATAPI-5 and SCSI-3 Block Commands provide a means to create a protected space on the secondary nonvolatile storage device 20a. As will be described below with reference to Fig. 5, the present invention combines these capabilities to create the BIOS protected area 27 on the secondary nonvolatile storage device 20a (such as in the hard disk drive 20a) to store a plurality of individual BIOS modules. Selected BIOS modules are read from the BIOS protected area 27 on the secondary nonvolatile storage device 20a into the system memory 13 and executed using the method 50 (Fig. 6) of the present invention to initialize the computer 10a.

The following is presented to better understand the PARTIES specification and its use in implementing the present invention. Protected Area Run-Time Interface Extensions Services (PARTIES) technology allows a system to reserve space on a hard drive for system use. This space is divided into segments via a Boot Engineering Extension Record (BEER). The individual segments can be used for data storage or for booting a fail-safe operating system.

When the system boots normally, the reserved space is inaccessible, and thus is protected from viruses, and unknowledgeable users, and the like. In one embodiment,

if a user elects to perform a fail-safe boot using MS-DOS, drive letters C: and above operate the way they normally would during a standard boot. The difference is that the system boots from drive A: instead of C:, where drive A: is simulated from one of the PARTIES segments. This capability has several applications including system
 5 diagnosis, recovery, and fail-safe applications.

In one implementation, there are four distinct software layers involved in using PARTIES. A Discovery layer detects the presence of PARTIES on the hard drive. A Boot Selection layer provides a user with the ability to choose the fail-safe boot service. A Simulation layer provides simulation services including drive A: from a reserved area
 10 on the hard drive when the user chooses to fail-safe boot the system. A Manipulation layer provides a way to create, delete and access PARTIES services. The ANSI PARTIES specification provides the specific details for formatting and finding PARTIES services.

In one implementation, during the discovery process, the BIOS checks all
 15 drives for the presence of a BEER sector. If a BEER sector is present, then the drive has PARTIES services available. If no BEER sectors are present then all PARTIES capability is disabled.

If fail-safe boot services are found during the discovery process, the user must be provided with methods to select a boot service. One implementation is as follows:
 20 (1) Integration with Phoenix Technologies MultiBoot 3 capability, and (2) Hotkey.

MultiBoot 3 provides two methods for boot selection. The first method is in SETUP, and the second method is via a hotkey during power on self test (POST). When PARTIES technology is present, the user is presented with a new option, Fail-Safe boot. When the user selects this option, a menu of boot services is displayed,
 25 using a MultiBoot 3 menu format.

In the case of a hotkey, the selected service is then booted. An OEM may wish to extend the system capabilities by providing buttons, or adding hotkeys, to trigger specific applications. Applications include compact disk (CD) control panels, System Diagnostics, and browsers. Another method of choosing a service to boot could
 30 involve a watchdog timer. If the primary OS fails to boot, the watchdog timer triggers and a failsafe boot is automatically performed. Still another method could involve adding buttons to the system. When a specific button is pressed the system automatically boots a specific failsafe service area. Still another possibility involves a touchpad or mouse. The user graphically selects a service area for the next boot.

35 After the user chooses to boot using PARTIES services, an INT 13 level simulation is invoked. A SETMAX command is issued to the drive that exposes the

entire service area, but leaves other service areas, further out on the drive, protected. Details of the simulation can be found in the ANSI PARTIES specification.

There are several possibilities for manipulating the PARTIES area. Firstly, a DOS-based tool can initialize the BEER sector as well as add and delete PARTIES
 5 services. Secondly, BIOS-based PARTIES services can be accessed from the BIOS SETUP utility as well as at run-time. The SETUP utility could allow the user to explicitly initialize create and delete BEER sectors and PARTIES services.

In a typical computer, the BIOS uses precious flash memory and CMOS
 10 memory to store non-volatile variables. In some BIOS implementations this can be 4k of memory or more. Furthermore, non-essential or non-differentiated BIOS modules such as those that are associated with operating system runtime functions such as BIOS Runtime Services and ACPI modules are also typically stored in either flash memory and CMOS memory. The present invention moves the storage location of these non-essential or non-differentiated BIOS modules to the secondary nonvolatile storage
 15 device 20a, such as the hard disk drive 20a, for example.

With reference to Fig. 4, it illustrates an exemplary layout of hard disk media 21 implemented in accordance with the principles of the present invention. Using the SETMAX command defined in the ATA/ATAPI-5 specification and the teachings of the PARTIES specification, a disk (media 21) layout illustrated in Fig. 4 may be created.

20 The hard disk media 21 of the hard disk drive 20a is configured in accordance with the principles of the present invention to have a PARTIES formatted area 27, a normal user area 26, and a vendor-protected area 25. The vendor-protected area 27 is protected from access by anyone but the manufacturer or vendor of the hard disk drive 20a. The normal user area 26 is used to store files and applications by the user in a
 25 normal fashion when using the computer 10a. The PARTIES formatted area 27, of BIOS protected area 27, stores a plurality of nonessential BIOS modules as was discussed above. These nonessential BIOS modules are read into system memory 11 as required by the dispatch manager 17 to set up the computer 10a for operation.

The BIOS retrieves information from the PARTIES formatted area 25 and uses
 30 the retrieved information to configure the system. The information stored in the PARTIES formatted area 25 can include option ROMs, BIOS utilities, or other data required to operate the computer 10a. In addition, the BIOS may use the PARTIES formatted area 27 to store variables in the same way that variables are stored in the critical nonvolatile storage device 15.

35 By way of example, in a computer system 10a with an ATA hard drive 20a as the secondary nonvolatile storage device 20a, for example, the procedure illustrated in Fig. 5 may be used to protect the BIOS modules. More particularly, Fig. 5 is a flow

diagram that illustrates an exemplary method 30 in accordance with the principles of the present invention that may be used to protect BIOS modules in a secondary nonvolatile storage device 20a. The method 30 comprises the following steps.

5 A secondary nonvolatile storage device 20a such as an ATA hard disk drive 20a, for example, is located 31. A READ NATIVE MAX command is issued 32 which returns the largest size of the ATA hard disk drive. A SETMAX command is issued 33 using the address returned by the READ NATIVE MAX command, which sets the size of the drive. A BEER sector (pointer) is located 34 using READ commands, which is in the last sector of the hard disk drive.

10 A PARTIES service area on the hard disk drive that contains system firmware is located 35. BIOS information is read from 36a and written to 36b the PARTIES service area as necessary. A SETMAX command is issued 37 to the SETMAX ADDRESS located in the BEER sector, which sets the size of the user area of the hard disk drive. A SETMAX PASSWORD command is issued 38, which is a random
15 number that changes every boot. A SETMAX LOCK command is issued 39 which allows no more SETMAX commands to be issued.

The normal operating system boot process is then initiated 40. The steps described above prevent unauthorized access to the PARTIES area by anything except the system firmware.

20 Fig. 6 is a flow diagram that illustrates an exemplary method 50 in accordance with the principles of the present invention that is used to store basic input output system (BIOS) modules of a computer 10a in a protected area, transfer selected BIOS modules to system memory 11 and execute the transferred BIOS modules. The method 50 comprises the following steps.

25 Minimal initialization code 16 corresponding to a small portion of the computer system BIOS is stored 51 in a first portion of a critical nonvolatile storage device 15. A dispatch manager 17 containing a plurality of tasks relating to configuring the computer 10a is stored 52 in a second portion of the critical nonvolatile storage device 15. A plurality of BIOS modules are stored 53 in a protected area of a secondary nonvolatile
30 storage device. The secondary nonvolatile storage device may comprise a hard disk drive, a compact disk ROM, a flash memory, a floppy disk drive, a Zip drive, or a SuperDisk drive, for example.

The computer 10a is turned on 54. The minimal initialization code is executed 55. The dispatch manager 17 is copied 56 from the critical nonvolatile storage device
35 15 into the system memory 11. The dispatch manager 17 is executed 57. The dispatch manager 17 sequences through the plurality of tasks relating to configuring the computer 10a to determine 58 if the task is stored in the system memory 11. If the task

is stored in the system memory 11, then the task is executed 59. The execute task step 59 may in fact load data and/or instructions from the secondary non-volatile storage device 20a as a result of executing the task. If the task is not stored in the system memory 11, the dispatch manager 17 finds 61 the code for the task and loads 62 the
5 code for the task in the system memory 11. Once the task is loaded in the system memory 11, the task is executed 59. The dispatch manager 17 iterates 64 through all tasks in the list of tasks until all tasks are loaded and executed. Once all tasks are loaded 63 the operating system of the computer 10a is launched 65.

Thus, methods for storing basic input output system (BIOS) modules of a
10 computer in a protected area, transferring the BIOS modules to memory, and executing the BIOS modules to configure the computer have been disclosed. It is to be understood that the above-described embodiments are merely illustrative of some of the many specific embodiments that represent applications of the principles of the present invention. Clearly, numerous and other arrangements can be readily devised by those
15 skilled in the art without departing from the scope of the invention.